

Method, system, device, signal and computer program product for metadata protection in
TV-Anytime

The invention relates to a method of providing data integrity authentication and data protection, in which a set of data fragments is protected by a signature.

The invention further relates to a system for providing data integrity authentication and data protection, the system being arranged to receive and handle data 5 fragments, of which a set of data fragments can be protected by a signature.

The invention further relates to a signature device for providing data integrity authentication and data protection, the device being arranged to handle data fragments, and the device being arranged to generate a signature to protect a set of data fragments.

The invention further relates to a verification device for verifying data 10 integrity authentication and data protection, the device being arranged to handle data fragments, and the device being arranged to verify a signature to protect a set of data fragments.

The invention further relates to a signal comprising data fragments, of which a set of data fragments is protected by a signature.

15 The invention further relates to a computer program product for implementing such a method.

As the number of channels available to television viewers has increased, along 20 with the diversity of the programming content available on such channels, it has become increasingly challenging for television viewers to identify television programs of interest. Historically, television viewers identify television programs of interest by analyzing printed television program guides. As the number of television programs has increased, it has become increasingly difficult to effectively identify desirable television programs using such 25 printed guides.

More recently, television program guides have become available in electronic format, often referred to as electronic program guides (EPGs). Like printed television program guides, EPGs present overviews of the available content, which can be browsed by the user. The general term content typically comprises things like music, songs, movies,

television programs, pictures and the likes, but can also refer to individual scenes, MPEG-4 objects, and so on.

The EPG compiles the overview from metadata that accompanies the individual content items. Metadata for content items is available from a variety of sources.

5 Metadata can be included with a broadcast stream (e.g. as MPEG-2 tables) or downloaded from external databases. For example, a television receiver or Personal Digital Recorder may be provided with an Internet connection, which allows the device to access metadata made available over the World Wide Web.

This metadata generally comprises information such as title, artist, genre and 10 so on, and may also contain a unique content reference identifier (CRID), sometimes also called a content reference identifier. Using the CRID, each individual content item can be uniquely identified. Further, using the CRID further information can be retrieved from a database. For example, a user can select a content item which he wishes to see from the EPG, even though the time and place of broadcast are not yet known. Using the CRID, the system 15 can then retrieve the time and place of broadcast of the content item when this information becomes available.

The CRID is not restricted to broadcast transmissions of content. It could also refer to a location on the Internet, or to any other source. The purpose of content resolution is to allow acquisition of a specific instance of a specific item of content. For example a user 20 may want to record an episode of a television series, but he does not necessarily know when and where that episode will become available. He can then use his personal digital recorder (PDR) or similar device to enter a reference to the episode or series by means of the CRID. Note that a CRID may refer to an entire series or to an individual episode thereof.

Having received a CRID for a content item, the PDR tries to obtain the 25 location of the content item. This information is called a locator and it contains the date, time and channel on which the content item will be broadcast. The user however does not need to be aware of this. Once the PDR has obtained the locator of the content item, the PDR waits for the specified date and time and then records the episode as it is broadcast on the specified channel. Of course, if the locator indicates a location on the Internet or the like, the PDR can 30 simply retrieve the content from the indicated location as soon as it becomes available.

The TV-Anytime standardization body provides a standardized Content Reference ID. See TV-Anytime Forum, www.tv-anytime.org, Specification Series: S-4, on Content Referencing (Normative), Document SP004V11, 14 April 2001, later version SP004V12, 28 June 2002, ETSI TS 102 822-4. This document specifies that the CRID

contains an <authority> field indicating the body that created the CRID. An authority will also provide the ability for the CRID to be resolved into locators or other CRIDs. A locator is the name for locations in time and space of content. The CRID further contains a <data> field which is a free format string that is compliant with the definition of Uniform Resource

5 Identifiers (URIs) as given in "RFC2396, Uniform Resource Identifiers (URI): Generic Syntax." This string should be meaningful to the authority given by the <authority> field.

The CRID is used for location resolution, which can be defined as the process of translating a CRID into other CRID(s) or locators. For instance, a CRID for an entire TV series could be translated into a series of CRIDs for the individual episodes of that series.

10 Location resolution may be done in the recording device (typically a Personal Digital Recorder or PDR) or remotely. A resolution provider does location resolution. Resolution providers use resolving authority records (RARs) to be identified and located. A RAR includes at least an <authority> field, corresponding to a body that creates CRIDs.

A RAR also contains a URL and the resolution provider name. The URL 15 points to the location where resolution information can be found. The resolution provider name contains the name of the body that is providing location resolution. These RARs are made available to PDRs.

TV-Anytime information and services are valuable so protection of this information is important. Protection includes the issue of source authentication and spoofing; 20 the integrity of the data is to be protected. When TVA data is received from a source, the receiver may want to check if the data is indeed coming from the expected source and hasn't been changed by a third party.

There is an incentive for a third party to try this. If a third party can change the metadata or CRID table, it can make the PDR record other information than was intended 25 including commercials, trailers or just other content. This is also very annoying for a user and may lower the trust the user has in the system. The PDR may therefore want to check whether the content came from a trusted source. If the data can be authenticated to originate from the source even when it is distributed using different channels, the PDR can use this to make a choice when confronted with multiple sources of the same content. An example of 30 this is when the data of a certain BBC show can be authenticated as being generated by the BBC, this raises the likelihood that this information is correct.

The source of the TV-Anytime data and metadata is not always the creator of the data. The source could be a service provider gathering and grouping information from different sources. It could be useful to check who created the data and whether the data has

been changed. In this case, the data that is received will hold parts provided by different sources.

The standard cryptographic approach to protection of data integrity is to sign the data using cryptographic techniques. As all TVA metadata is expressed in XML, a 5 transport neutral way of expressing signatures that allows the signatures to be carried in the same data structure would be to include the signatures in the TVA schema, and an obvious choice would be xmldsig ("RFC3275, (Extensible Markup Language) XML-Signature Syntax and Processing."). However, this standard uses the XPath data model ("XML Path Language (XPath) Version 1.0, W3C recommendation, J. Clark, S. DeRose, October 1999, 10 <http://www.w3.org/TR/1999/REC-xpath-19991116>") to define the evaluation of an XPath expression for every node of the XML parse tree, which is a transform which can be difficult to implement efficiently.

An attempt to overcome this problem has been described in "XML-signature XPath Filter" (W3C recommendation, latest revision 8 November 2002, 15 <http://www.w3.org/TR/xmldsig-filter>) which defines a XML signature transform to facilitate the development of efficient document subsetting. However, this recommendation is not adopted by the W3C consortium.

From this discussion, it is clear that there is an incentive for service providers and box manufacturers to use efficient integrity checking mechanisms for metadata.

20

It is an object of the invention to provide metadata integrity and source authentication to efficiently protect data fragments, enabling the protection of data fragments originating from different sources and enabling the protection of data fragments by multiple 25 authenticators. This will allow validation whether the data has not been changed during transport between (resolution) provider and the client. It will further allow validation whether the data has not been changed during the subsequent storage and handling of the data.

The object of the invention is achieved by a method according to the invention characterized in that each data fragment of the set comprises its own unique identifier, the 30 signature comprises references to the respective unique identifiers of the data fragments of the set. The invention describes a system in which signatures will be provided separately and references to unique identifiers will indicate which data fragments are covered by the signature. The at least one identifier uniquely identifies a data fragment in order to enable the linking between data fragments and the signature. This is done by providing an (optional)

field that is added to each data fragment for identification of that data fragment. An existing field could be used if all data fragments have such a field that uniquely identifies each data fragment. Otherwise, a special signature identifier could be added to each data fragment as an optional field. When this field is present it is a unique identification of that data fragment

5 instance within the data. A signature may refer to multiple data fragments, such that sets of data fragments are signed instead of individual data fragments. This has the further advantage that it is more efficient.

An embodiment of the method according to the invention is described in claim 2. During distribution, when data changes hands, more than one party may apply their 10 signature to the same data fragment. These signatures may apply to different subsets of data fragments, i.e., totally disjunct subsets, partially overlapping subsets, or equal subsets.

An embodiment of the method according to the invention is described in claim 3. An advantage is that only the hashes are needed to compute or verify a signature. This is especially advantageous (reduction of computation time) if the same data fragments 15 are being used in multiple signatures.

An embodiment of the method according to the invention is described in claim 4. XML clearly separates individual data fragments enabling the identification of data fragments in a standardized way. Data fragments can be collected in one or more XML documents.

20 An embodiment of the method according to the invention is described in claim 5. As described before, TV-Anytime metadata requires protection from unauthorized manipulating of metadata. The invention can therefore be advantageously be applied in the TV-Anytime environment.

An embodiment of the method according to the invention is described in 25 claim 6. A suitable choice for data expressed in XML would be to use the xmldsig definition of a signature, with the references to the unique identifiers added.

An embodiment of the method according to the invention is described in claim 7. Data fragments to be signed can be approached by the use of a transform function (according to RFC3275) that removes the data fragments that are not considered for this 30 signature. The transform function refers to the data fragments using the unique identifier.

An embodiment of the method according to the invention is described in claim 8. As is also explained in the aforementioned xmldsig specification, the same text can be coded in multiple ways using different encodings. In order to calculate the signature, one defined representation of the document has to be defined. This process is called

cannolization. Indication which cannolization function is used, allows the calculation of the signature values without the need of extracting the data first.

An embodiment of the method according to the invention is described in claim 9. In order to protect the integrity of the references, the references themselves might 5 either implicitly or explicitly be included in the data to be signed.

An embodiment of the method according to the invention is described in claim 10. More elaborate search options could be provided by adding signature index files.

Such an index file would then link references to the appropriate signature files using the unique identifier to support a search option. This table provides a grouping between 10 the data that is signed and the list of signatures.

An embodiment of the method according to the invention is described in claim 11. In order to ensure that the identifier is unique among data fragments of the same type within this instance of data it is suggested to start the unique identifier with a unique identification of the organization responsible for generating the data fragment. This would 15 also allow the client to detect what organization published the data.

An embodiment of the method according to the invention is described in claim 12. The DNS name is an easy and understandable choice for the unique identification of the organization.

An embodiment of the method according to the invention is described in 20 claim 13. Although the unique identifier identifies a data fragment, it does not define where this data fragment can be found within the total data. In order to facilitate the searching of the correct data fragment within the data, the reference should preferably also be accompanied by a location indicator.

An embodiment of the method according to the invention is described in 25 claim 14. An implementation indicates the path through the data that has to be taken in order to locate the data fragment.

An embodiment of the method according to the invention is described in claim 15. It is possible and efficient to include the signature information within the data document.

An embodiment of the method according to the invention is described in 30 claim 16. A different approach is to define a wrapper around the data document that further includes the signature information and possibly some other elements that need signing. In this way, the original data is included in the wrapper without changes, except for the addition of

possibly missing unique identifiers. A suitably defined wrapper may be extended to allow additional data to be included in the signed data.

An embodiment of the method according to the invention is described in claim 17. A different approach is to define a separate data document comprising the signature information, referring to the data fragments in the original data document. In this way, the original data document remains unchanged, except for the addition of possibly missing unique identifiers.

The system according to the invention is characterized in that the system comprises means to receive and handle data fragments of the subset, each data fragment identified by a unique identifier, the system further comprises at least one of -means for associating a signature with the protected data fragments of the set using their unique identifiers, -means for verifying a signature associated with the set using the unique identifiers of the protected data fragments, and -means for generating a signature that references the protected data fragments by their unique identifiers.

The signature device according to the invention is characterized in that the device is arranged to address each data fragment to be protected by a unique identifier included in the data fragment, and the device is arranged to generate signature information comprising the unique identifiers to refer to the data fragments of the set.

The verification device according to the invention is characterized in that the device is arranged to address each data fragment to be protected by a unique identifier included in the data fragment, and the device is arranged to verify signature information comprising the unique identifiers to refer to the data fragments of the set.

The signal according to the invention is characterized in that each data fragment of the set comprises its own unique identifier, and the signature comprises references to the unique identifiers of the data fragments of the set.

The invention is further characterized by a computer program product for implementing the method of claim 1.

These and other aspects of the invention will be further described by way of example and with reference to the schematic drawings, wherein:

Fig. 1 schematically illustrates the process of content resolution,
Fig. 2 shows the different fragments as defined by TV-Anytime,

Fig. 3a, 3b, and 3c show XML definitions and examples related to the unique identifier,

Fig. 4a, 4b, and 4c show XML definitions and examples related to the signature information, and

5 Fig. 5 shows XML definitions and examples related to the key information.

As an illustration of the collection and handling of data, handling of metadata will be described by a device such as a personal digital recorder or PDR in the TV-Anytime 10 environment. Fig. 1 schematically illustrates the process of content resolution. The PDR 10 is instructed to record a content item identified by a Content Reference Identifier CRID. Instructing the PDR to record a content item, or in other words scheduling that content item for recording can be done in a variety of ways. A presently common way is that the user 15 manually indicates, e.g. by selecting the content item in the EPG, that the content item is to be recorded. It will be readily understood that part or all of the functionality ascribed to the PDR below could also be incorporated into one or more other devices, such as television receivers, set-top boxes or personal computers. The computer program product 12 with the computer-readable instructions in a suitable format such as a optical disc or solid state 20 memory can be used to store or distribute the program instructions implementing the invention.

The PDR, or another device to which the PDR is connected, may be equipped to determine kinds of content items that the consumer may be interested in. This is known as user profiling or recommender systems. By keeping track of content items which the consumer views, and employing an implicit and/or explicit rating system for such content 25 items, it becomes possible to predict with varying degrees of accuracy which other content items the consumer may be interested in. It then becomes possible to automatically record content items which are likely to be of interest to the consumer. Such content items could then be recorded by the PDR. Many techniques for user profiling are known in the art. When the PDR determines, using user profiling, that a particular content item may be of interest, it 30 schedules the content item for recording.

The CRID for the content item is used to facilitate automatic recording of the content item. The CRID could be entered manually by the user, or be the result of selecting a content item through an Electronic Program Guide. This second option assumes that the CRID is somehow provided to the PDR together with other metadata used in the EPG by a

CRID provider entity 13. Alternatively, if the CRID is not known by the user or by the PDR, the user could perform a search using for example the title of the content item in a metadata database, and select the desired content item from the search results. The CRID is then supplied to the PDR by the search engine.

5 There are many other ways to provide the CRID to the PDR. For example, a trailer or preview for a movie could be broadcast with the CRID embedded in the content of the commercial in some way (e.g. a watermark). The user could then press a button on his remote control, television or PDR. The PDR or television then extracts the CRID from the content of the commercial.

10 Once the CRID for the desired content item is known, the PDR tries to obtain locator information for the content item, using the CRID as input. This locator information is not necessarily always available. For example, the CRID may refer to a movie that has only recently been released in movie theaters. This movie is not likely to be broadcast on television in the near future, so it cannot be scheduled using EPG information. In such a case, 15 the PDR should regularly try to obtain the locator, as the locator may become available later (e.g. a year later, when the movie is going to be broadcast on TV). The CRID could also refer to a TV series, which is then resolved into a number of CRIDs for individual episodes of that series. It is possible that no locator information is available for some episodes. Here the PDR should also regularly retry to obtain the locator(s) for those episodes.

20 The process of translating a CRID into locator information is known in TV-Anytime as location resolution. Location resolution involves mapping a location-independent content reference (the CRID) to its location in time (e.g. scheduled transmission time in a broadcast system) and space (e.g. TV channel, IP address). As explained above, these locations in time and space are referred to as "locators." The process of location resolution 25 may happen inside the PDR or by using a physically remote server, such as a server on the Internet.

To the PDR, the CRID essentially contains opaque information, which it cannot resolve to a location without external assistance. A Resolution Provider (RP) which provides locator information for CRIDs is provided to solve this problem. Usually multiple 30 RPs are available, and the PDR must know which RP to use for a particular CRID. Often, this is the same body that created the CRID. The name of the authority is present in the CRID in the <authority> field, as explained above. This name is present in the form of a registered Internet domain name. It is possible for an Resolution Authority (RA) 15 to be found on the

Internet using the domain name resolution process specified in the TV-Anytime specification SP004.

Each RA will require one or more Resolving Authority Records (RAR) to exist in the PDR for location resolution to take place. Each resolving authority record will 5 need to be placed inside some sort of transport specific container which allows the PDR to know that this is a RAR. In the case of multiple records for the same authority, the PDR can choose to just use one of them, or try them all in turn. The Resolving Authority Record (RAR) contains the information that identifies the RAs where content reference resolution information can be found.

Using the RAR, the PDR determines which RP to use to resolve a particular CRID. The PDR then submits a request for a location accompanied by a CRID to the Resolution Provider in question. In response to this request, the Resolution Provider returns the locator information (assuming this information is available in that RP, of course). The PDR can then access the content source and obtain the content item. A content item may 15 have more than one locator, for example if it is broadcast multiple times or available from multiple providers. The PDR may then choose which locator to use, or prompt the user to make a selection.

Once the locator information has been obtained, the PDR waits for the specified date and time and then records the episode as it is broadcast on the specified 20 channel. Of course, if the locator indicates a location on the Internet or the like, the PDR can simply retrieve the content from the indicated location as soon as it becomes available.

Content items for which locator information is available can be recorded by the PDR at the appropriate moment. To this end, the PDR may comprise local storage such as a sufficiently large hard disk, and/or a device such as a DVD+RW writer. The storage on 25 which content items are stored needs not be local to the PDR, but may also be an external device such as a hard disk or a file server connected to the PDR via a home network. Once the content items have been recorded, they can be played back at any time until they have been erased.

Using the above approach, anyone knowing the location of content could act 30 as a resolution provider. Content and service providers, however, may desire that only authorized resolution providers perform content resolution for their content, for example to be able to protect their reputation. On the other hand, for consumers and PDRs it is important to be able to rely on/trust the CRID authority and resolution provider, so that they can obtain the correct content.

If the PDR operates in accordance with a Digital Rights Management (DRM) system, then a content item may be erased when the rights associated with the content item require such erasure. Also, some content items may not come with a right to record the item at all, or with a right that permits viewing only for a limited amount of time, or for a limited number of times. The PDR should then erase the content item when the limit is exceeded, or refuse further access to the content until further rights are obtained that permit further access. The content as received in the client box can be protected during transport by encryption. Before the content can be accessed, the content has to be decrypted. This process is controlled by the DRM or conditional access (CA) system.

10 The TV-Anytime specification distinguishes between two different distribution methods: unidirectional and bidirectional. In the unidirectional situation, TV-Anytime data is another stream in the broadcast stream with the normal signaling in place. Access to this stream can be protected using traditional conditional access systems, such as scrambling. Using the normal signaling methods defined for the transport mechanism, the 15 conditional access system is identified and the messages carrying the conditional access information related to this stream are indicated. Most digital broadcast systems use the MPEG-2 transport stream format (ISO/IEC 13818-1:1996(E), Information technology - Generic coding of moving pictures and associated audio information: Systems, First Edition, 1996-04-15).

20 In the bidirectional case, a point to point connection is made between client and server. This process is described in TV-Anytime document SP004v1.2, Specification series S-4 on Content Referencing, Version 1.2, Final Specification, 28 June 2002, ETSI TS 108 822-4. The DRM system will open a secure channel to the service provider and tunnel the communication through this channel. In this way it will ensure that only authorized TV-25 Anytime clients can access the content.

Although existing CRC mechanisms in the broadcast system will deal with transmission errors, it remains wishful to detect intentional changes and to authenticate that the information was generated by the claimed source.

30 The source of the TV-Anytime data is not always the creator of the data. The source could be a service provider gathering and grouping information from different sources. Data can also be retrieved by the PDR from different sources. Thus, data that is received will hold parts provided by different sources. It could be useful to check who created to data and whether the data has been changed. This is done using signatures.

All TVA metadata is provided as TVA fragments. In TVA, according to the metadata specification (TV-Anytime document WD647/SP003v1.3 Part A, Specification series S-3 on Metadata: Part A Metadata Schemas, version 1.3, 15 December 2002, ETSI TS 102 822-3), a TVA fragment is "a self contained atomic portion of the metadata". In this 5 document, it is assumed that the smallest TVA metadata element that can be signed is a fragment.

The invention includes the definition of a label that is used to uniquely identify a TV-Anytime fragment in order to link the fragment to the signature. This is done by providing an optional field that is added to each TV-Anytime fragment. When present, this 10 unique identifier is a unique identification of that fragment instance within this instance of metadata. The unique identifier should allow for easy tracing of the fragment within the metadata. This is required in order to find the different fragments that are needed to calculate a signature.

Fig. 2 shows the different fragments as defined by TV-Anytime. In order to 15 support signatures, all fragments have an optional or compulsory unique identifier for identification of a fragment. Within the TV-Anytime specification, a field called TVAID is used in these fragments. According to the above-identified metadata specification TVAIDs are used to "indicate uniqueness within a metadata description". Although they seem to match the requirement for an identifier, they're only unique for a particular type of TVAID. 20 e.g. a serviceID and segmentID could be the same within a particular metadata description.

Several variations to implement the concept of a unique identifier are possible.

In a first variation of the invention, the TVAID is used as reference identifier. This provides a unique identifier if the reference used in the signature indicates the context (e.g. service or segment). The TVAID can be used if all fragments have one (or one is added) 25 and it is determined that using the TVAID a unique reference to the fragment can be made within this instance of the TVA metadata.

In a second variation, a special TVA signature identifier is added to all fragments as an optional field. Either the TVAID or a new identifier is defined for this purpose. The XML definition for an example identifier added to each TVA fragment that 30 uniquely identifies a fragment among other fragments of the same type within the metadata (or within the relevant instance of the metadata) is shown in Fig. 3a. In order to be able to reference each fragment (or set of fragments), an example of a format for all TV-Anytime fragments could be to add a TVASignatureId attribute, formally defined as shown in Fig. 3b.

A further advantageous variation of the invention ensures that the identifier is unique among fragments of the same type within the metadata by starting the identifier with the DNS name (or another unique identification) of the organization responsible of generating the fragment. This would also allow the client to detect what organization 5 published the data. An example TVASignatureId of a fragment published by company MyCompany could look as shown in Fig. 3c.

Labeling each individual fragment or set of fragments and signing using references has the advantage that, when properly chosen, the reference would provide a link 10 from the signature file to the fragment. Furthermore, the unique identifier provides a link between the fragment(s) and the data containing the signature.

As all TVA metadata is expressed in XML, a transport neutral way of expressing signatures that allows the signatures to be carried in the same data structure could be to include the signatures in the TVA schema. As TV-Anytime metadata is expressed in 15 XML a suitable choice would be xmldsig, but other XML-based signature schemes can of course be defined as well.

The signature can be stored according to the xmldsig standard. Additionally, references to the unique identifier indicate which fragments are used to calculate a signature. Although the unique identifier identifies a fragment, it does not define where this fragment 20 can be found within the total TVAMain. In order to facilitate the searching of the correct fragment within the metadata, the reference (defined in URI format according to RFC2396) should preferably also indicate the location. An extension of the invention therefore adds an optional location indicator to indicate the path through the metadata that has to be taken in order to locate the fragment.

A compliant example definition of the fragment URI including the unique 25 identifier according to the invention can be defined as "tva://<path>/<TVASignatureId>" where <path> is the path from the start of the metadata towards the fragment, and TVAID is the identifier of the fragment.

Some examples are "tva://TVAMain/aap.org;132423", "tva://TVAMain/ClassificationTable/CSAlias/publisher.com;122314", and 30 "tva://TVAMain/ProgramDescription/ProgramLocationTable/Schedule/metwt.org;320984". The first example illustrates that it is also possible to sign the complete TVAMain metadata document. In this case, all of TVAMain without the certificate and signature parts should be considered.

Because in this embodiment the unique identifier (such as TVASignatureId) is used in URIs to refer to a fragment, the identifier should be compatible with the formatting restrictions placed upon URIs as described in RFC2396. Furthermore, in order to ease the parsing of the URI no slashes ("/") are used in the TVASignatureId.

5 The TVASignatureTable will provide a grouping between the data that is signed and the list of signatures; an example definition of such a table is shown in Fig. 4a and Fig. 4b. In this definition, there is an option to include other TV-anytime metadata documents such as the ContentReferencingTable and ResolvingAuthorityRecordTable in this table. This has the advantage that, as they can only occur once, no unique identifier is needed for that 10 metadata in the table and in the URI.

A TVASignatureTable is defined as shown in Fig. 4c. Zero or more signatures can be present. As data may not, not yet, or not anymore be available in a system, not all signatures available for the data indicated in TVASignatureWrapper are always included and not all fragments that are protected by the various signatures are always present. The 15 implementation of the delivery system will have to take care that the relevant fragments and signature are present when needed. In the bi-directional delivery system missing fragments or missing signatures can be downloaded.

A different way to define the fragments that need to be signed is by defining a transform function (according to RFC3275) that removes some or all elements from the 20 metadata that are not considered for a signature.

In order to check a signature, the corresponding public key of the party that applied the signature(s) is needed. Distribution of these keys can be done in several ways. They can be hard coded in the device, but as this would raise problems if new keys are used or when current keys are compromised, the most common way of distribution of the keys is 25 by incorporating them into a so-called certificate-chain ("Applied Cryptography Second Edition: protocols, algorithms, and source code in C, Bruce Schneier, Wiley, 1996"). So in order to check the signature, in addition to the signatures data, also the certificates of the parties providing signatures are required. TVASignature allows the inclusion of one or more ds:KeyInfo objects in order to support the carriage of such certificates within the 30 TVASignature wrapper. A XML complex type indicating a list of KeyInfo objects with accompanying identifiers is shown in Fig. 5.

In order to be able to refer from a signature to KeyInfo elements in the KeyInfoWrapperTable from the Signatures, the reference URI is defined as "tva://KeyInfoListTable/<Identifier>", where <Identifier> is the identifier indicated in the

KeyInfoWrapperType. Some examples are therefore "tva:// KeyInfoListTable/132423", "tva:// KeyInfoListTable/435432h", and "tva:// KeyInfoListTable/MyKeyInfo".

This allows the inclusion of certificates and other options of communicating KeyInfo objects. It also indicates how they are linked to signatures.

5 (Public) keys can also be stored using X509 certificates, where the subject field in the 509 certificate may too contain the organization name or a different unique identification identifying the key source. This provides an additional way to make sure the data is really signed by the organization that claims to have signed it.

10 Signatures can be generated using a suitable algorithm, such as DSA or RSA signature generation algorithms.

15 As is explained in the xmldsig specification, text can be coded in many ways. In order to calculate the signature, one defined representation of the document has to be defined. This process is called cananolization. Within TV-Anytime, BiM is used as a binary codex for the binary encoding of TV-Anytime data. When BiM encoding is used, BiM should be indicated as the cananolization function. This allows the client to use the BiM encoded files to calculate the signature values without the need of extracting the data first. The transform and cananolization function are used to process the content so a unique, always the same, set of bytes is produced over which a signature is calculated.

20 In order to also protect the integrity of the references to the unique identifier, they can either implicitly or explicitly be included in the data to be signed. As described above and relating to Fig. 5, signatures can also be used to protect different tables used in a system.

25 More elaborate search options could be provided by adding signature index files. Such an index file would than link unique identifiers to the appropriate signature files.

30 As is explained in "ISO/IEC 13818-6:1998, Information technology -Generic coding of moving pictures and associated audio information: Extensions for Digital Storage Media Command and Control, 1998." and "ETSI TS 102 812 V1.1.1 (2001-11), Digital Video Broadcasting (DVB); Multimedia Home Platform (MHP) Specification 1.1, 28 June 2001." Digital signatures can be implemented by calculating a hash over the content and signing the hash using public key cryptography.

There are two variations of the invention with respect to the location of the signatures.

In a first variation of the invention, the metadata object (such as TVAMain) is expanded with the signature fragments so it will include the signature information. This is

advantageous if the signatures are distributed and accessed within the normal distribution system as is indicated by TV-Anytime.

In a second variation of the invention signatures are provided separately, for example by definition of a wrapper that includes the TVAMain and possibly some other elements that need signing. This is advantageous as this would not change the current 5 metadata specification and it would also allow to include other TV-Anytime documents (e.g. ContentReferencingTable and ResolvingAuthorityRecordTable).

The system according to the invention supports signatures over single or multiple fragments. This has the advantage that it enables to sign a set of fragments, which is 10 more efficient than signing each individual fragment. It has the further advantage that it remains as much as possible compatible with the unidirectional as well as the bi-directional distribution system, while minimizing the level of change upon the existing metadata specification.

The measures used in the above embodiments can be used individually, but 15 these measures could also be combined to provide for better protection, or for protection against multiple threats.

It should be noted that the above-mentioned embodiments illustrate rather than limit the invention, and that those skilled in the art will be able to design many alternative embodiments without departing from the scope of the appended claims.

20 In the claims, any reference signs placed between parentheses shall not be construed as limiting the claim. The word "comprising" does not exclude the presence of elements or steps other than those listed in a claim. The word "a" or "an" preceding an element does not exclude the presence of a plurality of such elements. The invention can be implemented by means of hardware comprising several distinct elements, and by means of a 25 suitably programmed computer.

In the device claim enumerating several means, several of these means can be embodied by one and the same item of hardware. The mere fact that certain measures are recited in mutually different dependent claims does not indicate that a combination of these measures cannot be used to advantage.

30 Of course, the techniques above can also be used both inside and outside the scope of TV anytime.